

RISK-INFORMED DECISION MAKING Application to Technology Development Alternative Selection

Homayoon Dezfuli ⁽¹⁾, Gaspare Maggio ⁽²⁾, Christopher Everett ⁽²⁾

⁽¹⁾ Office of Safety & Mission Assurance, NASA Headquarters, 300 E Street SW, Washington, D.C., USA
Email: hdezfuli@nasa.gov

⁽²⁾ Technology Risk Management Operations, ISL, 350 Seventh Avenue, Suite 401, New York, NY, USA
Email: gmaggio@islinc.com and ceverett@islinc.com

ABSTRACT

NASA NPR 8000.4A, Agency Risk Management Procedural Requirements, defines risk management in terms of two complementary processes: Risk-Informed Decision Making (RIDM) and Continuous Risk Management (CRM). The RIDM process is used to inform decision making by emphasizing proper use of risk analysis to make decisions that impact all mission execution domains (e.g., safety, technical, cost, and schedule) for program/projects and mission support organizations. The RIDM process supports the selection of an alternative prior to program commitment. The CRM process is used to manage risk associated with the implementation of the selected alternative. The two processes work together to foster proactive risk management at NASA. The Office of Safety and Mission Assurance at NASA Headquarters has developed a technical handbook to provide guidance for implementing the RIDM process in the context of NASA risk management and systems engineering. This paper summarizes the key concepts and procedures of the RIDM process as presented in the handbook, and also illustrates how the RIDM process can be applied to the selection of technology investments as NASA's new technology development programs are initiated.

1. INTRODUCTION

Risk management (RM) is an integral aspect of virtually every challenging human endeavor, but the complex concepts that RM encapsulates and the many forms it can take makes it difficult to effectively implement. However, few will disagree that effective risk management is critical to program and project success.

Until recently, NASA's RM approach has been based on Continuous Risk Management (CRM), which stresses the management of individual risk issues during implementation. In December of 2008, NASA revised its RM approach, in order to more effectively foster proactive risk management. This approach, which is outlined in NPR 8000.4A, *Agency Risk Management Procedural Requirements* [1], evolves NASA's risk management to entail two complementary processes: Risk-Informed Decision Making (RIDM)

and Continuous Risk Management (CRM). The former is intended to inform decisions through better use of risk and uncertainty information in selecting alternatives. The latter is used to manage risks over the course of the development and implementation cycle to assure that safety, technical, cost, and schedule goals are met. Before, RM was considered equivalent to the CRM process; now, RM is defined as comprising both CRM and RIDM processes, which work together to assure proactive risk management as NASA programs and projects are conceived, developed, and executed. Figure 1 illustrates the concept.




Figure 1. Risk Management as the Integration of Risk-Informed Decision Making (RIDM) and Continuous Risk Management (CRM)

This paper addresses the RIDM component of RM, as discussed in the NASA Risk-Informed Decision Making Handbook [2]. In some respects this could be considered the more important part of the overall RM activities, since the decisions made during program initiation ultimately "burns in" the risk that must be retired/mitigated during the life of the program (primarily during the development portion of the life cycle) using the CRM process to track and sustain progress towards the program's objectives.

In addition, although the RIDM process articulated in the RIDM Handbook is focused on programs and projects, it is also easily adaptable to decision making in other contexts. Indeed, future versions of the RIDM Handbook will expand the guidance to specifically address other decision contexts, such as institutional decision making. However, given NASA's current proposed shift in emphasis towards the development of space exploration enabling technologies, there is an immediate need for guidance on the application of RIDM in support of technology development. Consequently, this paper aims not only to illustrate the RIDM process as set forth in the RIDM Handbook, but

also to show how the process can be used to risk-inform the selection of an effective technology investment strategy.

2. THE RIDM PROCESS

As specified in NPR 8000.4A, the RIDM process itself consists of the three parts shown in Figure 2.

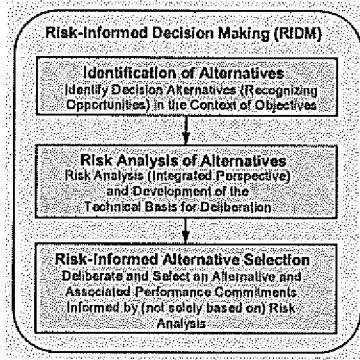


Figure 2. The RIDM Process

2.1. Part 1, Identification of Alternatives

2.1.1. Deriving Performance Measures from Objectives

Objectives, which in general may be multifaceted and qualitative, are captured through interactions with the relevant stakeholders. They are then decomposed into their constituent derived objectives, each of which reflects an individual issue that is significant to some or all of the stakeholders. At the lowest level of decomposition are quantifiable performance objectives, each of which is associated with a performance measure that quantifies the degree to which the performance objective is met. Typically, each performance measure has a “direction of goodness” that indicates the direction of increasing benefit.

A comprehensive set of performance measures is considered in decision making, reflecting stakeholder interests and spanning the mission execution domains

of *Safety* (e.g., avoidance of injury, fatality, or destruction of key assets), *Technical* (e.g., meeting payload mass requirements), *Cost* (e.g., execution within allocated cost), and *Schedule* (e.g., meeting milestones). A notional objectives hierarchy is shown in Figure 3.

Objectives that have natural unit scales (e.g., Minimize Cost, Maximize Payload) are generally easy to associate with appropriate performance measures (e.g., Total Cost [\$], Payload Mass [kg]). Other objectives might not have an obvious or practical natural unit scale, requiring the development of constructed scale or proxy performance measures.

Performance measures that use constructed scales are typically appropriate for measuring objectives that are essentially subjective in character. An example of such an objective might be Maximize Stakeholder Support. Here, there is no natural measurement scale by which an assessment of stakeholder support can be made. Instead, a subjective scale (e.g., 1 to 5) can be constructed that supports subjective quantification of stakeholder support.

Alternatively, proxy performance measures can be used to indirectly quantify an objective without recourse to subjective assessment. Properly chosen proxy performance measures provide a quantifiable metric that correlates with (as opposed to corresponds to) the objective of interest. For example, the number of attendees at monthly meetings might be a satisfactory proxy performance measure for stakeholder support. Figure 4 shows the relationship between natural, constructed scale, and proxy performance measures.

Objectives whose performance measure values must remain within defined limits give rise to imposed constraints that reflect those limits. Imposed constraints limit the set of compiled decision alternatives to those that can potentially meet them, thus warranting the investment of resources required for further analysis.

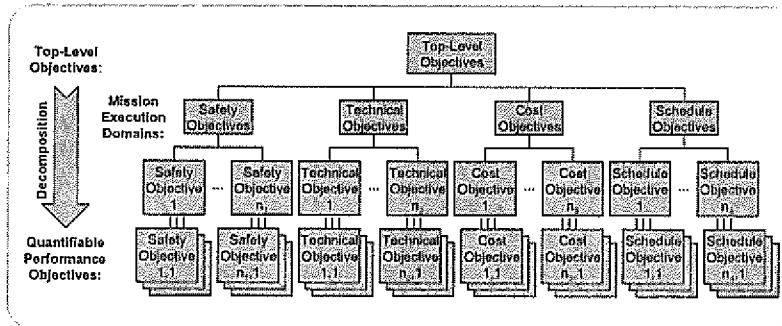


Figure 3. Notional Objectives Hierarchy




Figure 4. Types of Performance Measures

2.1.2. Compiling Feasible Alternatives

One way to generate decision alternatives under consideration is by a trade tree [3]. Initially, the trade tree contains high level decision alternatives representing high level differences in the strategies used to address objectives. The tree is then developed in greater detail by determining general categories of options that are applicable to each strategy. Trade tree development continues iteratively until the leaves of the tree contain alternatives that are well defined enough to allow quantitative evaluation via risk analysis. Along the way, branches of the trade tree containing unattractive categories are pruned, as it becomes evident that the alternatives contained therein are either infeasible or categorically inferior to alternatives on other branches.

2.1.3. Technology Investment - Identification of Alternatives

In the context of technology investment, Part 1 of the RIDM process begins with a vision of the level of capability that investment in the development of enabling technologies is intended to achieve. In his address at the Kennedy Space Center on April 15, 2010, President Obama stated, "Our goal is the capacity for people to work and learn and operate and live safely beyond the Earth for extended periods of time, ultimately in ways that are more sustainable and even indefinite." This top-level vision must be decomposed into a set of essential mission capabilities that together represent the overall capability implied by the goal. Once these essential capabilities are defined, they serve as the basis for identifying the specific functional capabilities required by each mission. The ability to attain these functional capabilities becomes the yardstick by which the efficacy of a proposed technology investment strategy can be measured.

2.1.3.1. Technology Investment - Deriving Performance Measures

Like top-level mission objectives, top-level objectives that are stated as capabilities may also be multifaceted and qualitative, and need to be decomposed into

specific capabilities whose attainability can be quantitatively assessed. As indicated above, this is done by decomposing the top-level capabilities into mission capabilities. Each mission capability should be unique in that its capability is not implied by some combination of other mission capabilities. Cumulatively, the mission capabilities should fully address the top-level capability objectives.

The objectives hierarchy is further refined by specifying, for each mission capability, one (or more) *mission objectives* that bound the levels of performance (e.g., payload capacity, stay duration, crew size) needed to enable the capability in general. Figure 5 notionally illustrates a simplified "agency-level" objectives hierarchy for technology investment. Each mission objective then serves as the "top-level" objective as depicted in Figure 3 such that safety, technical, cost, and schedule related performance measures are defined for each. In this way, the performance measures for each mission milestone would serve as a measure of the incremental functionality each technology development provides with time towards the ultimate goal of 100% functionality (100% being defined as the highest degree of functionality that would be required to enable all of the mission milestones). However, unlike mission-oriented programs, which have specific cost and schedule constraints, these performance measures remain unconstrained for the purposes of making risk-informed technology investment decisions. The reason for this is that the true measure of performance will be the return on the investment on a set of technologies (a.k.a. a *technology suite*), i.e., how long will it take a given technology development strategy to enable each mission objective (assuming a fixed budget profile for all technology suite alternatives).

2.1.3.2. Technology Investment - Compiling Feasible Alternatives

It is the nature of technological innovation to focus on the development of functionality, often in the absence of a specific application and without cognizance of other development activities that may have beneficial




Figure 5. Mission Objectives Hierarchy (Technology Investment)

synergies (or incompatibilities) in certain contexts. As applied to technology investment, the RIDM process ties technology development to mission objectives by first conducting a functional decomposition of each mission objective, identifying the functions that must be performed by any solution that is ultimately implemented. The functional decomposition facilitates the identification of safety and technical performance measures. A high degree of functional overlap is expected among the mission objectives, although with varying degrees of performance depending on mission specifics. Nevertheless, the functional decomposition produces a set of functions around which to propose, through brainstorming and other means, a set of candidate technologies from which different *technology suites* can be identified.

Figure 6 shows a notional technology/functionality mapping matrix that associates each proposed technology (rows) to the function(s) it performs (columns).

	Function 1	Function 2	Function 3	Function 4	Function 5	Function 6	Function 7	Function 8	Function 9	Function 10	Function 11	Function 12	Function 13	Function 14	Function 15	Function 16	Function 17
Technology 1																	
Technology 2																	
Technology 3	x			x													
Technology 4		x															
Technology 5																	
Technology 6																	
Technology 7																	
Technology 8																	
Technology 9																	
Technology 10																	

Figure 6: Technology/Functionality Mapping matrix

At a minimum, a feasible technology suite must collectively address all of the functions. Note that all feasible technology suites must contain Technologies 1, 3, 6, 7, & 8 (if any of these technologies fails to materialize then one or more functions would be absent and therefore one or more milestones could not be achieved). A trade tree (see Figure 7) can be used to generate the set of all feasible technology suites. Any

technology suite that is incapable of providing ALL of the necessary functions is deemed infeasible (these are indicated by RED highlighting in Figure 7). Additionally, any technology suite that depends on incompatible technologies is also infeasible (e.g., no technology suite in Figure 7 contains both Technology 4 and Technology 10). The alternatives (technology suites) that remain following pruning are forwarded for risk analysis.

2.2. Part 2, Risk Analysis of Alternatives

2.2.1. Program/Project-Oriented Risk Analysis of Alternatives

For each feasible alternative, the performance measures are quantified, taking into account whatever significant uncertainties stand between the decision to implement the alternative and the accomplishment of the objectives that drive the decision making process to begin with. Given the presence of uncertainty, the actual outcome of a particular decision alternative will be only one of a spectrum of outcomes that could result from its selection, depending on the occurrence, nonoccurrence, or quality of occurrence of intervening events. Therefore, it is incumbent on risk analysts to model each significant possible outcome, accounting for its probability of occurrence, to produce a distribution of forecasted outcomes for each alternative, as characterized by probability density functions (pdfs) over the performance measures (see Figure 8 on the next page).

Each mission execution domain quantifies the performance measures under its purview according to domain-specific processes and methods, such as those in the NASA Cost Estimating Handbook [4] in the case of cost. For the purpose of RIDM, two aspects of risk analysis are paramount as outlined on the next page.

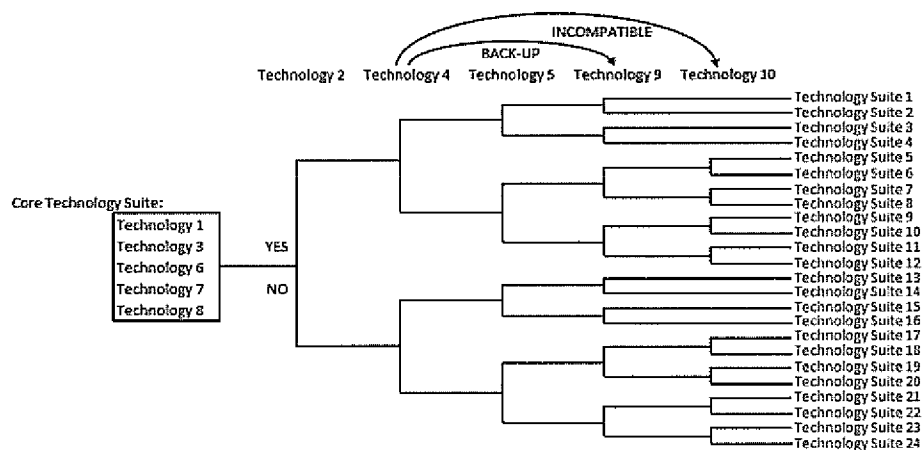


Figure 7. Notional Technology Suite Alternative Trade Tree

- Risk analysts conduct RIDM at a level sufficient to support robust selection of a preferred decision alternative. If the uncertainty of one or more performance measures is preventing the decision maker from confidently assessing important differences between alternatives, then the risk analysis is iterated to provide additional uncertainty reduction where it is relevant to decision making. The analysis stops when the technical case is made; if simpler, more qualitative methods are sufficient to produce a robust decision, then more detailed methods need not be applied. In other words, a graded approach to risk analysis is applied, driven by the need for robust decision making.
- Domain-specific analyses should be integrated into a multidisciplinary risk analysis framework that: operates on a common set of *performance parameters* (with associated uncertainties) across alternatives and mission execution domains (e.g., the cost model uses the same mass data as the lift capacity model); consistently addresses uncertainties across alternatives and mission execution domains (e.g., budget uncertainties, meteorological variability); and preserves correlations between performance measures. Figure 9 shows a notional risk analysis framework which, for a given alternative, transforms a defining set of (uncertain) performance parameters into a set of (uncertain) performance measures. In order to preserve correlations, the risk analyses in this framework should be exercised using consistent sampling of uncertain performance parameters. One way to accomplish this is for the entire framework

to be exercised within a single common Monte Carlo shell.

The principal product of Part 2 of the RIDM process is the Technical Basis for Deliberation (TBfD), a document that: catalogues the objectives hierarchy and the set of candidate alternatives; summarizes the analysis methodologies used to quantify the performance measures; and presents the risk analysis results. The TBfD is the input that risk-informs the deliberations that support decision making. The presence of this information does not necessarily mean that a decision is risk informed; rather, without such information, a decision cannot be said to be risk informed.

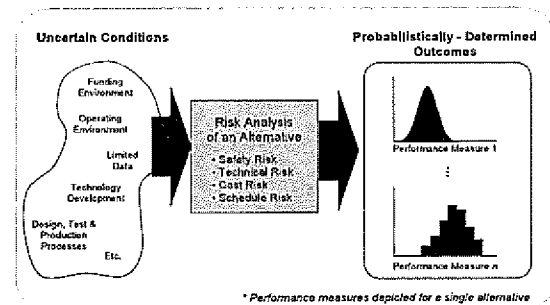


Figure 8: Uncertainty of Forecasted Outcomes Due to Uncertainty of Analyzed Conditions

2.2.2. Technology Investment - Risk Analysis of Alternatives

In the case of technology development decision

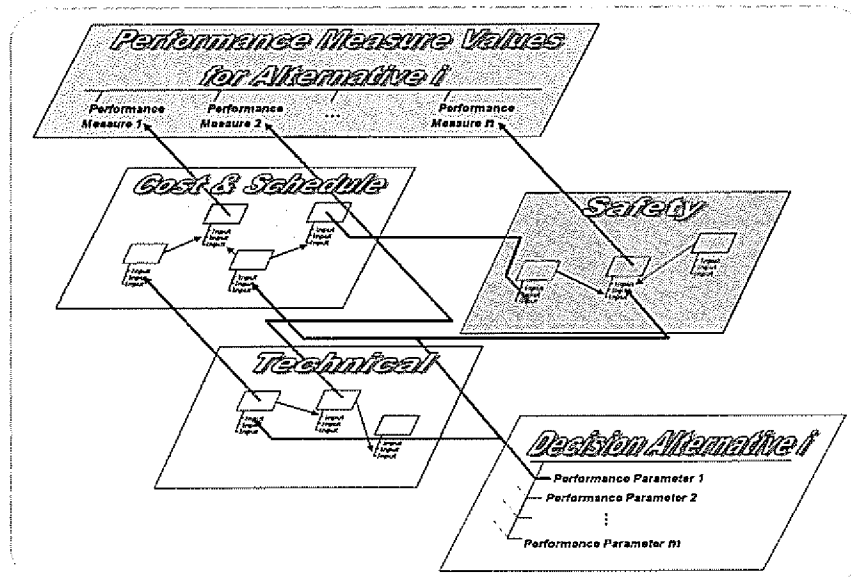


Figure 9: Risk Analysis Framework (Alternative Specific)

making, the performance of each technology suite is assessed in terms of the technology development plan for each alternative technology suite. Each plan introduces (uncertain) costs and schedules into the risk assessment based on an analysis of what it takes to develop the underlying technologies in the context of the overall technology suite.

The plan for a single technology suite alternative is shown notionally in the milestone/functionality matrix of Figure 10. The matrix shows a series of specific milestones that are set to gauge progress towards the objective of 100% functionality across all identified functions. The percentages listed in the matrix of Figure 10 notionally indicate that different milestones may involve different levels of performance from the functions. Thus, the order of the milestones is closely related to the technology development plan, which is specific to each technology suite. The milestones themselves may be technology development milestones at the program level or missions selected at the agency level from among the mission objectives. Regardless of the specific nature of the milestones, the analysis must indicate when each mission objective has been enabled by the development of the involved functions at sufficient levels of performance to meet mission needs. Different alternatives will, in general, contain different milestones and/or different orderings of shared milestones, but all alternatives share a common goal, namely the enabling of the mission objectives.

	Function 1	Function 2	Function 3	Function 4	Function 5	Function 6	Function 7	Function 8	Function 9	Function 10	Function 11	Function 12	Function 13	Function 14	Function 15	Function 16	Function 17
Milestone 1	10%			10%		40%				30%							
Milestone 2	40%		30%	10%			70%					50%			20%	20%	
Milestone 3	90%	10%	10%		40%	50%	10%		70%		70%		10%	10%			
Milestone 4		40%		10%		70%			30%								
Milestone 5			40%	70%			70%			40%							
Milestone 6						70%					40%				40%		
Milestone 7		50%				80%	40%		80%		60%	80%	60%				40%
Milestone 8			50%		80%	80%	50%			70%	60%		80%	60%	50%		
Milestone 9	60%	80%		100%	100%	100%	100%	100%	40%	100%		50%	100%	100%	100%	100%	50%
Milestone 10	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%

Figure 10: Milestone/Functionality Mapping matrix

2.3. Part 3, Risk Informed Alternative Selection

2.3.1. Performance Commitments

Performance measure pdfs constitute the fundamental risk analysis results needed to risk-inform subsequent deliberations and decision making. However, there are practical difficulties comparing performance measures whose values are expressed as pdfs, since for one thing, they represent ranges of potentially overlapping values. In many decision contexts it is appropriate to compare alternatives in terms of their expected performance. However, in the presence of performance thresholds, over-reliance on expected performance in decision making has the potential to:

- Introduce potentially significant probabilities of falling short of imposed constraints, thereby putting

objectives at risk, even when the mean value meets the imposed constraints; and

- Contribute to the development of derived requirements that have a significant probability of not being achievable.

In order to circumvent these issues, the RIDM Handbook introduces the concept of performance commitments. A performance commitment is a performance measure value set at a particular percentile of the performance measure's pdf, so as to anchor the decision maker's perspective to that value as if it would be his/her commitment, were he/she to select that alternative. For a given performance measure, the performance commitment is set at the same percentile for every decision alternative, so that the probability of failing to meet it is the same across alternatives, even though the performance commitments themselves differ from one alternative to the next. Performance commitments are not themselves performance requirements. Rather, performance commitments are used to risk-inform the development of credible performance requirements as part of the overall systems engineering process.

The use of performance commitments in RIDM supports a risk-normalized comparison of decision alternatives, in that a uniform level of risk tolerance is established prior to deliberating the merits and drawbacks of the various alternatives. Put another way, risk normalized performance commitments show what each alternative is capable of, at an equal likelihood of achieving that capability, given the state of knowledge at the time.

The inputs to performance commitment development are:

- The performance measure pdfs for each decision alternative (or, more generally, the joint distribution over the performance measures);
- An ordering of the performance measures; and
- A risk tolerance for each performance measure, expressed as a percentile value.

For each alternative, performance commitments are established by sequentially determining, based on the performance measure ordering, the value that corresponds to the stated risk tolerance, conditional on meeting previously defined performance commitments. This value becomes the performance commitment for the current performance measure, and the process is repeated until all performance commitments have been established for all performance measures.

Once the performance commitments are developed, each alternative can be compared to every other alternative in terms of their performance commitments, with the deliberators understanding that the risk of not achieving the levels of performance given by the performance commitments is the same across alternatives.

2.3.2. Technology Investment - Risk Informed Alternative Selection

In the case of technology development alternative selection, for each technology suite, the risk analysis produces a pdf for each mission objective, which characterizes when the mission objective will be enabled by the technology development plan. Performance commitment development then consists of the decision maker applying his/her mission-objective-specific risk tolerances to the pdfs to produce risk-normalized *schedule commitments* for the attainment of each mission capability. Figure 11 notionally presents how three different technology suites may perform at enabling mission objectives for three notional mission milestones. The domain of every performance measure is time, with the direction of goodness being sooner rather than later. Additionally, there may be hard schedule constraints that mark some long-duration development plans as infeasible (i.e. manned asteroid mission by 2025). These infeasible plans would be selected out as part of the deliberation process.

2.3.3. Deliberation

The risk informed alternative selection process within

RIDM provides a method for integrating risk information into a deliberative process for decision making, relying on the judgment of the decision makers to make a risk informed decision. The decision maker does not necessarily base his or her selection of a decision alternative solely on the results of the risk analysis. Rather, the risk analysis is just one input to the process, in recognition of the fact that it may not model everything of importance to the stakeholders. Deliberation employs critical thinking skills to the collective consideration of risk information, along with other issues of import to the stakeholders and the decision maker, to support decision making.

Deliberations are structured to include the following activities:

- *Comparing performance commitments:* Comparing alternatives in terms of their performance commitments facilitates deliberation and decision making because it separates issues of performance from issues of achievability (i.e., risk). The deliberation establishes performance commitments based on assessments of their own (or the decision maker's) risk tolerance. Performance commitments may shift during deliberations as the deliberators influence one another concerning risk tolerance and performance expectation; they are only finalized when the decision maker selects an alternative for implementation.
- *Additional uncertainty considerations:* Although performance commitments provide a scalar means of assessing performance that incorporates the deliberators' and decision maker's risk tolerances,




Figure 11: Performance Commitments and Risk Tolerances for Three Alternatives